AFRL-IF-RS-TR-2001-203
Final Technical Report
October 2001

ADAPTIVE CONTROL OF DISTRIBUTED AGENTS THROUGH PHEROMONE TECHNIQUES AND INTERACTIVE VISUALIZATION

Environmental Research Institute of Michigan

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

20020116 202

ADAPTIVE CONTROL OF DISTRIBUTED AGENTS THROUGH PHEROMONE TECHNIQUES AND INTERACTIVE VISUALIZATION

H. Van Dyke Parunak
Jeff Posdamer
John A. Sauter
Sven Brueckner

REPORT DOCUMENTATION PAGE

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	Oct 01	Final Aug 99 - Mar 01

4. TITLE AND SUBTITLE
ADAPTIVE CONTROL OF DISTRIBUTED AGENTS THROUGH PHEROMONE TECHNIQUES AND INTERACTIVE VISUALIZATION

5. FUNDING NUMBERS
C   - F30602-99-C-0202
PE - 63760E
PR - J116
TA - 00
WU - 01

6. AUTHOR(S)

H. Van Dyke Parunak, Jeff Posdamer, John A. Sauter and Sven Brueckner

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Environmental Research Institute of Michigan

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Defense Advanced Research Projects Agency
3701 North Fairfax Drive
Arlington, VA 22203-1714

AFRL/IFSA
26 Electronic Pky
Rome, NY 13441-4514

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

AFRL-IF-RS-TR-2001-203

11. SUPPLEMENTARY NOTES

AFRL Project Engineer: Carl A. Defranco, IFSA, 315-330-3096

12a. DISTRIBUTION AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
Software agents guided by synthetic pheromones can imitate the stigmergetic dynamics of insects. The resulting software architecture is well suited to problems such as the control of air campaigns and unmanned robotic vehicles. An architecture for the deployment of synthetic pheromones in a military setting is described. The architecture can be instantiated into a single computer, a distributed network of computers, or even across a mass network of field processors implanted in the environment. Software agents representing the military units of interest move through this synthetic environment sensing and depositing synthetic pheromones. These agents are able to locate and select targets dynamically. Using ant colony optimization techniques, the agents build safe paths to the target that adapt in real time to new threats or opportunities. The results produced can be used to help guide manned missions, or to directly control unmanned robotic vehicles. The technology has particular advantages for controlling large swarms of small, unmanned vehicles.

14. SUBJECT TERMS

Fine-grained agents, stigmergy, pheromones, biomimetics, air campaign control

15. NUMBER OF PAGES
32

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

SUMMARY

Software agents guided by synthetic pheromones can imitate the stigmergetic dynamics of insects. The resulting software architecture is well suited to problems such as the control of unmanned robotic vehicles. We introduce the approach, describe the mechanisms we have developed, and summarize the technology's performance in a series of scenarios reflecting military command and control.

1. INTRODUCTION

Most recently, in DARPA's JFACC program [11], we have applied this architecture to the problem of controlling air combat missions, with special emphasis on unmanned air vehicles. In the course of our experimentation, we have developed several mechanisms that are promising for agent coordination in general. This report describes pheromone-based movement control as a variety of potential-field-based methods (Section 2), reviews the mechanisms we have developed (Section 3), describes their performance in several air combat scenarios (Section 4), and describes how the architecture could be implemented (Section 5).

2. POTENTIAL FIELDS VIA PHEROMONES

From an engineering perspective, pheromones are a particularly attractive way to construct a potential field that can guide coordinated physical movement.

2.1 Potential Fields

Many social insect species coordinate the activities of individuals in the colony without direct communication or complex reasoning. Instead, they deposit and sense chemical markers called "pheromones" in a shared physical environment that participates actively in the system's dynamics. The resulting coordination is robust and adaptive. Seeking such characteristics in engineered systems, we have developed a software runtime environment that uses synthetic pheromones (data structures inspired by the insect model) to coordinate computational agents using mechanisms similar to those of social insects.

Potential-based movement systems are inspired by electrostatics. The (vector) electric field $\vec{E}(\vec{r})$ at a point in space is defined as the force felt by a unit charge at that point. We define a (scalar) potential field $\varphi_{21} = -\int_{P_1}^{P_2} \vec{E} \bullet d\vec{r}$ by integrating this vector field from an arbitrary reference point to each point in the space. Conversely, the field may be expressed as the gradient of the potential, $\vec{E} = -\nabla \varphi$, and a massless charged particle will move through space along this gradient.

In electrostatics, the field is generated by the physical distribution of charges, and may be computed by Coulomb's law. Einstein's extension of the formalism to gravity leads to a gravitational field generated by the physical distribution of mass. Thus the movement of a massive charged particle will follow a composition of two fields.

The notion of movement guided by a potential gradient has been applied to other situations in which the field is generated, not by natural physical phenomena, but by synthetic constructs. A parade example is robot navigation [14], which automatically maps from a given distribution of targets and obstacles to a movement plan. In such applications, the designer of the field is not limited to two components of the field (electrostatic and gravitational), but can include many different fields to represent different classes of targets and obstacles.

We are interested in using a potential field to guide unmanned robotic vehicles (URV's) through the battlespace (Figure 1). In this scenario, robotic vehicles seek to destroy the tank farm, which is defended by two missile batteries. The vehicles climb a potential gradient centered on the tank farm while avoiding gradients centered on the threats. To be useful in warfighting, this field requires four characteristics (mnemonically, "4-D"):

Diverse.—It must fuse information of various types and from various sources, including targets to be approached, threats to be avoided, and the presence of other URV's with whom coordination is required.

Distributed.—Centralized processing of a potential field imposes bottlenecks in communications and processing, and generates localized vulnerabilities to attack. Ideally, the potential field should be stored close to where the information that it integrates is generated, and close to where it will be used.

Decentralized.—Efficiency and robustness also dictate that components of the system be able to make local decisions without requiring centralized control, ideally on the basis of nearest-neighbor interactions with one another.

Dynamic.—The battlespace is an uncertain and rapidly changing environment, and the methods and architecture used to construct and maintain the field must be able to incorporate such changes rapidly into the field.

Physical Assets   **Field Representation**

Figure 1: Potential Fields Corresponding to Physical Assets

An architecture inspired by insect pheromones satisfies these requirements, and can be applied effectively to warfighting scenarios.
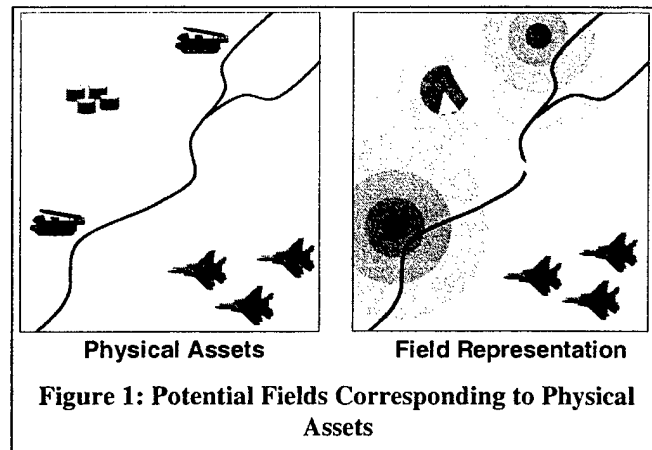
2.2 Synthetic Pheromones

Insects perform impressive feats of coordination without direct inter-agent coordination, by sensing and depositing pheromones (chemical scent markers) in the environment [10]. For example, ants construct networks of paths that connect their nests with available food sources. Mathematically, these networks form minimum spanning trees [5], minimizing the energy ants expend in bringing food into the nest. Graph theory offers algorithms for computing minimum spanning trees, but ants do not use conventional algorithms. Instead, this globally optimal structure emerges as individual ants wander, preferentially following food pheromones and dropping nest pheromones if they are not holding food, and following nest pheromones while dropping food pheromones if they are holding food.

The real world provides three operations on chemical pheromones that support purposive insect actions.

- It *aggregates* deposits from individual agents, providing fusion of information across multiple agents and through time.

- It *evaporates* pheromones over time. This dynamic is an innovative alternative to traditional truth maintenance. Traditionally, knowledge bases remember everything they are told unless they have a reason to forget something, and expend large amounts of computation in the NP-complete problem of reviewing their holdings to detect inconsistencies that result from changes in the domain being modeled. Ants immediately begin to forget everything they learn, unless it is continually reinforced. Thus inconsistencies automatically remove themselves within a known period.

- It *diffuses* pheromones to nearby places, disseminating information for access by nearby agents.

The pheromone field constructed by the ants in the environment is in fact a potential field that guides their movements. Unlike many potential fields used in conventional robotics applications, it satisfies the 4-D characteristics:

Diverse.—Ants can respond to combinations of pheromones, thus modifying their reaction to multiple inputs at the same time.

Distributed.—The potential field is generated by pheromone deposits that are stored throughout the environment. These deposits do their work close to where they are generated, and are used primarily by ants that are near them.

Decentralized.—Both ant behavior and pheromone field maintenance are decentralized. Ants interact only with the pheromones in their immediate vicinity, by making deposits and reading the local strength of the pheromone field. Because diffusion falls off rapidly with distance, deposits contribute to the field only in their immediate vicinity.

2

Dynamic.—Under continuous reinforcement, the pheromone field strength stabilizes rapidly, as a concave function of time (proportional to $\int_0^t E^\tau d\tau$ where $E \in (0,1)$ is the evaporation rate) [2]. Thus new information is quickly integrated into the field, while obsolete information is automatically forgotten, through pheromone evaporation.

An implementation of synthetic pheromones has two components: the *environment* (which maintains the pheromone field and performs aggregation, evaporation, and diffusion), and the *actors* (which deposit and/or react to the field maintained by the environment). Our implementation has two corresponding species of agents. A set of *place agents* with a Neighbor relation defining adjacency makes up the environment, and each actor is represented by a *avatar agent*. In Hindu mythology, the term refers to an incarnation of a deity; hence, an embodiment or manifestation of an idea or greater reality. In our system, an avatar is the manifestation in our system of the greater reality (ground truth in the battlespace). An avatar travels with the physical entity that it represents. It moves from one place agent to another only when its parent entity moves physically from one region to another. Later we will introduce a third species of agents called *ghosts*. Table 1 summarizes the three species of agents in ADAPTIV.

Table 1: Species of Agents in ADAPTIV

Agent	Description	Number
Avatar	Represents a physical entity (e.g., plane, mission package, URV) in the battlespace	1 per physical entity
Place	Represents a region in the battlespace	1 per region
Ghost	Does stochastic search for an avatar	Many per avatar

Each place agent maintains a scalar variable corresponding to each pheromone flavor. It augments this variable when it receives additional pheromones of the same flavor (whether by deposit from an avatar or by propagation from a neighboring place), evaporates the variable over time, and propagates pheromones of the same flavor to neighboring place agents based on the current strength of the pheromone. The underlying mathematics of the field developed by such a network of places, including critical stability theorems, rest on two fundamental equations [2]. The parameters in both are:

- $P = \{p_i\}$ = set of places

- $N: P \to P$ = neighbor relation between places. Thus the places form an asymmetric multigraph.

- $s(t,p)$ = pheromone strength at time t and place p

- $r(t,p)$ = external input at time t to place p

- $q(t,p)$ = propagated input at time t to place p

- $E \in (0,1)$ = evaporation parameter

- $F \in [0,1)$ = propagation parameter

The first equation describes the evolution of the strength of a single pheromone flavor at a given place.

$$s(t+1,p) = E * s(t,p) + r(t,p) + q(t,p)$$

The first term of this equation reflects evaporation of pheromone, the second reflects external deposit of new pheromone, and the third reflects propagation from neighboring places. That propagation is in turn described by the second fundamental equation:

$$q(t+1,p) = \sum_{p \in N(p')} \frac{F}{|N(p')|} \left(r(t,p') + q(t,p')\right)$$

3

Using these equations, one can demonstrate several critical stability and convergence theorems, including:

- *Local Stability*: The strength of the output propagated from any set of places to their neighbors at $t + 1$ is strictly less than the strength of the aggregate input (external plus propagated) to those places at t.

- *Propagated Stability*: There exists a fixed upper limit to the aggregated sum of all propagated inputs at an arbitrary place if a one-time and one-place external input is assumed.

- *Global Stability*: The pheromone strength in any place is bounded.

Our implementations include another parameter, the pheromone *threshold S*. If the strength of the pheromone at a location drops below this threshold, the software no longer processes that pheromone, and it disappears from the system.

In principle, there are no restrictions on the graph of place agents. In physical movement problems, each place agent is responsible for a region of physical space, and the graph of place agents represents adjacency among these regions. There are different ways in which place agents can be assigned to space. In JFACC, we tile the physical space with hexagons, each representing a place agent with six neighbors.

An avatar agent is associated with one place agent at any given time. It can read the current strength of pheromones at that place as a function of their flavors, and deposit pheromones into the place. It can also determine from the place agent the relative strength of a given flavor at the place and at each of its neighbors. An avatar moves from one place to another by spinning a roulette wheel whose segments are weighted according to this set of strengths.

Such techniques can play chess [4] and do combinatorial optimization [1], and we have applied them to manufacturing [2] and military C^2 [11].

3. BASIC MECHANISMS

We have explored several basic mechanisms essential to the engineering deployment of pheromone mechanisms. These fall into four broad categories: combinations of multiple pheromones, including history in movement decisions, ghost agents, and visualization. Some of the results discussed in this section are expounded at more length in other publications, but are drawn together here so that they can be more readily considered as an integrated system.

3.1 Pheromone Vocabulary

There are two ways in which the pheromone vocabulary can be multiplied. First, different flavors may reflect different features of the environment (e.g., Red (hostile) air defenses, Blue (friendly) bases). These flavors have different *semantics*. Second, different flavors with the same semantics (e.g., all generated by the same feature) may differ in their evaporation or propagation rate or threshold, thus having different *dynamics*.

Pheromones with Different Semantics.—We explored the effect of increasing the semantics of a pheromone vocabulary in the context of a pheromone adaptation of AI's classic missionary-cannibal problem [12]. Three missionaries and three cannibals find themselves together on one bank of a raging river, with a dugout canoe capable of carrying only one or two people. If at any time the cannibals outnumber the missionaries on either bank of the river, they will eat them. The problem is to plan a sequence of moves that gets all six people safely across the river.

At each decision epoch, only those agents on the bank with the boat make a movement decision. Each such agent decides whether it wishes to move by evaluating a personal choice function that returns a real number between 0 and 1, evaluating a random variable uniformly distributed on [0,1], and comparing these two values. If the random number is less than the value of the choice function, the agent places itself on a list of candidates for movement. The actual riders in the boat are chosen randomly from the list of candidates.

The interesting details of the agent's decision are embedded in its choice function, which is a function of the levels of the available pheromones. In principle, each individual agent could have its own choice function, but in our experiments all Missionaries share one choice function and all Cannibals share another.

We explore the performance of the system for various combinations of three distinct pheromones: a bank pheromone that tells agents where they are, an undifferentiated population pheromone deposited by both Missionaries and Cannibals, and distinctive Missionary and Cannibal pheromones. Our performance metric is the number of steps necessary for the system to move the agents from one bank to the other. Because of the stochastic nature of the decisions, different runs often yield different numbers of steps, and we report the median run length over 100 runs.

Figure 2 shows the result for one series of experiments, comparing three different pheromone configurations. "R" indicates the performance for agents executing a blind random walk. "BSmart" shows the performance available when the agents have access only to a pheromone indicating which bank they inhabit (thus one bit of information). The performance at "BPSmart" results from telling them in addition the total population on their bank. Since there are six possible populations on either of

Figure 2: Performance and Pheromone Vocabulary (Error bars show inter-quartile spread)

two banks, the information available is $\log_2(2*6) = 3.58$. "BMCSmart" reflects the performance when missionaries and cannibals deposit distinct pheromones. There are $4*4$ possible equilibrium values on each bank, but no agent will ever sense the combination $\{0,0\}$, so the total information available is $\text{Log}_2(2*(4*4-1)) = 4.91$. Figure 2 shows that log performance is linear in information content, so performance is exponential in information content.

In the missionary-cannibal experiments, the agent's choice function explicitly takes into account the levels of the different pheromones. An alternative approach, used in our air combat applications, computes a weighted function of the various input pheromones to create a single "net pheromone" whose gradient avatars then follow. In this case, the basic pheromone flavors are:

- RTarget: emitted by a red (hostile) target.

- GTarget: emitted by a blue (friendly) agent who has encountered a red target and is returning to base.

- GNest: emitted by a blue agent who has left the base and is seeking a target.

- RThreat: emitted by a red threat (e.g., Surface-to-Air Missile)

In addition, we provide the blue agent with Dist, an estimate of how far away the target is.

Initially, we experimented with an equation of the form

$$\frac{6 \cdot RTarget + \gamma \cdot GTarget + \beta}{\alpha \cdot RThreat + \delta \cdot Dist + \beta}$$

where α, β, γ, δ, and θ are tuning factors, easily manipulated in a genetic algorithm or particle swarm optimization. β avoids singularities when other terms are 0. This form attracts blue agents to targets or to the trails of other blue agents who have found targets, avoids threats, and seeks to minimize distance to the target. While yielding reasonable performance, this equation left some performance gaps. Manual manipulation of the equation yielded the alternative form

$$\frac{6 \cdot RTarget + \gamma \cdot GTarget + \beta}{(\rho \cdot GNest + \beta)(Dist + \varphi)^{(\delta + \alpha(RThreat + 1))} + \beta},$$

which gives much improved performance. While superficially more complex, this latter equation could readily be discovered by genetic mechanisms.

5

Pheromones with Different Dynamics.—Another technique involving multiple pheromones is to use pheromones with the same semantics but differing dynamics (e.g., rates of evaporation E and propagation F and threshold S) [3]. To motivate this mechanism, consider the distribution of pheromone sources shown in Figure 3. Each source (or background 0) is at one cell of a hexagonal grid.

We are interested in the degree of guidance that the pheromone field offers an avatar located at a given place. Let f_i be the pheromone strength at place i. The guidance g_j available to an avatar at place j is

$$ g_j = \underset{i \in \{j\} \cup N(j)}{Max} \left(\frac{f_i}{\sum f_i} \right) - \frac{1}{(1+N(j))} $$

Guidance thus ranges from 0 (if all accessible places have the same pheromone strength) to 1 (if only one place has pheromone and all the others have none).

Figure 3: Test Distribution of Pheromone Sources

Figure 4 shows the distribution of guidance (white = 1, black = 0) for two different propagation parameters F. When F is low (left plot), most places in the target-rich region at the left of the figure have high guidance, but the pheromones do not propagate across the targetless right side of the figure, yielding a broad "valley" with low guidance. When F is high (right plot), propagation merges signals from individual sources, yielding low guidance in the target-rich region but a much narrower valley on the right. Thus high propagation gives good long-range guidance but poor short-range guidance, while low propagation gives good short-range guidance but poor long-range guidance.

Figure 4: Guidance fields for low (left) and high (right) propagation parameters

A reasonable resolution is to have each source deposit multiple pheromones with different dynamics. An avatar picks its next step first by measuring the guidance available from each flavor, then computing its movement based on the pheromone with highest guidance. Figure 5 shows the guidance field from six flavors with different dynamics, yielding both high guidance in the target area, and propagation of pheromones across most of the eastern valley.

3.2 History

The movement of an avatar through the graph of places should balance several factors. A strong field gradient enables deterministic hill climbing that the avatar should exploit. However, a weak gradient may result from noise in the system. In


Figure 5: Guidance field with multiple pheromones

6

this case, it does not provide reliable guidance. We would prefer that the avatar continue moving in the general direction of its previous steps if there is one, and otherwise that it explore more broadly.

To balance deterministic hill climbing and stochastic exploration, the avatar moves from one place to another by spinning a roulette wheel whose segments are weighted according to the relative strengths of a pheromone flavor (or weighted combination of flavors) in the place and its neighbors. The mapping function from relative pheromone weight to segment width determines the degree of stochasticity in the avatar's behavior. For example, if $s_i{}'$ is the perceived pheromone concentration at place i, we compute a normalized weight p_i for that place as:

$$p_i = \frac{s_i{}'}{\sum_{j=1} s_j{}'}$$

where the summation ranges over place i and its neighbors, and then compute the probability p_i' that the avatar will move to that place as:

$$p_i' = \frac{e^{\beta * p_i}}{\sum_{j=1} e^{\beta * p_j}}$$

The parameter β determines the degree of stochasticity in the avatar's movement. On a hex grid, when $\beta < 4$, selection probabilities are more similar than the pheromone strengths would indicate, favoring exploration, while $\beta > 5$ tends to emphasize stronger gradients, favoring exploitation.

To balance hill climbing against previous direction, we assign momentum to the avatar. Models of actual ant behavior usually restrict the ant's ability to smell pheromones to some angle on either side of its current orientation. In our implementation, this technique takes the form of multiplying each segment in the avatar's roulette wheel by a weight that is strongest in the direction the avatar is currently heading, and weakest in the direction from which it has just come.

Such a momentum works well if the avatar is moving over continuous space. However, representing (continuous) space as a (discrete) graph of place agents can introduce anisotropies that confuse a simple momentum computation. For example, Figure 6 shows five geodesics on a hex lattice. Trajectories a, b, and c maintain a constant heading, but trajectories d and e experience local direction changes while executing a shortest path across the lattice. A straightforward momentum function will interfere undesirably with these necessary changes of direction. To avoid this problem, each avatar maintains an exponentially-weighted moving average of its past headings and modulates the relative strengths of the pheromones in its vicinity by a measure of the angular alignment between each candidate place and the current value of the heading history.


Figure 6: Path anisotropy in a hex lattice

7

3.3 Ghost Agents

Of particular interest to robotic applications is the emergence of discrete paths in the pheromone field as a large population of ants concurrently read and reinforce it. For example, Figure 7a shows the pheromone field deposited by a swarm of ants wandering out from their nest (at the lower left of the figure) in search of food (at the upper right). Initially, the field is roughly circularly symmetrical, and serves to guide food-bearing ants back home. Once some ants find the food and begin returning home, this field rapidly collapses into a path (Figure 7b).


a. Initial *Field*   b. Final *Path*

Figure 7: Path Condensation

At first glance, this dynamic [6] violates second-law tendencies to increasing disorder in systems consisting of many components. Left to themselves, large populations tend to disorder, not organization. Natural systems can organize themselves at the macro level because their actions are coupled to a flow field at a micro level. Agents *perceive* and orient themselves to the flow field and reinforce that field by their *rational action*, as shown by the solid lines in Figure 8 [8]. Metaphorically, they drain unwanted entropy from the macro level (where organization is desired) to the micro level (where disorder is tolerated).

Traditional coordination mechanisms ignore the micro level completely, as agents perceive and act directly on one another (dashed line in Figure 8). We link

Figure 8: Coordination through Environmental Fields

agents through the environment so that perception and action serve both to coordinate multiple agents and to control overall disorder.

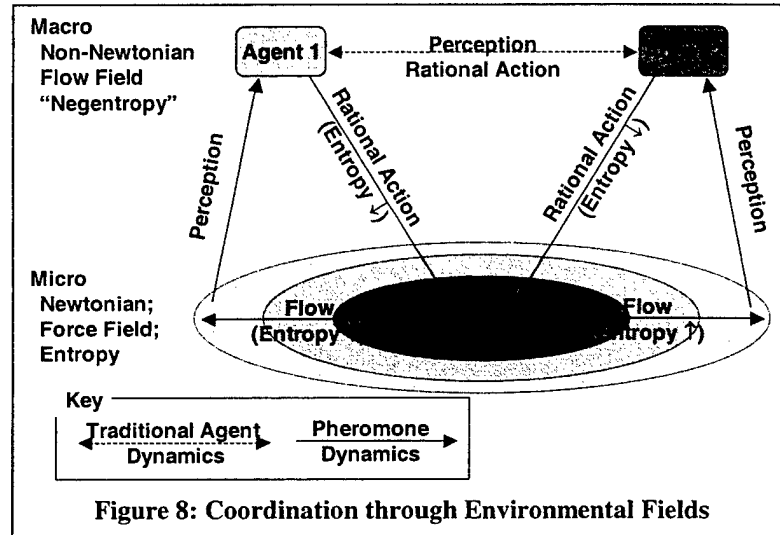We validate this mechanism of emergent coordination explicitly through experiments that compute the entropy over time of the pheromone molecules at the micro level and the agents at the macro level [13]. The increase in entropy at the micro level (through Brownian motion of pheromone molecules) more than balances the decrease in entropy experienced by avatars following the pheromone gradient.

The path emergence illustrated in Figure 7 is the result of interactions among many avatars. Each avatar's behavior is highly stochastic, performing a real-time Monte Carlo search of its local vicinity, and contributing to the emergence of a long-range path. In engineering applications, it may not be feasible to ask hundreds of physical robots to explore the domain in this manner, nor is it necessary.

An avatar may send out many unembodied representatives, which we call *ghosts*. Ghosts can move as fast as the network among place agents can carry them. As they sense pheromones in the environment, they deposit their own flavors of pheromones, which condense into paths like that in Figure 7. The avatars then follow this path. As an avatar moves, it continuously sends out ghosts, so that the path it follows is being constantly revised to accommodate dynamic changes in the environment.

Our experiments show this path formation dynamic to be extremely robust and adaptive. Figure 9 shows the formation of a path from a friendly airbase (lower right) to the nearer (in terms of safest path) of two targets (the

8

Figure 9: Path to the nearer of two targets.

Figure 10: Threading a gauntlet.

house-shaped icons), avoiding threats (the radar icons). If we increase the strength of the left-hand target to twice that of the closer target, the path will lead there instead. In all cases the ghosts were able to identify a path of minimum distance that led to the target while avoiding air threats.

Figure 10 shows the formation of a path to a target protected by a long narrow gauntlet of threats oriented away from the base. This is a particularly challenging problem since the ghosts must determine that a safe path exists even though it involves moving far way from the target. This particular configuration resisted classical potential field methods. These methods required the identification of an intermediate waypoint near the mouth of the gauntlet. When the ghosts were presented with the same configuration except that the end of the gauntlet was sealed, they settled on a path that pierced the air defense at the target end, thus identifying the closest, safest path to the target.

The presence of multiple targets provides another opportunity for ghosts: selecting the closest and most important target. In our infrastructure, the importance (or priority) of a target is encoded in its strength, which defines how much pheromone it deposits into the environment. There is some trade-off that needs to be made between distance and strength of the target. We have explored the balance between these factors by setting up two targets T1 and T2 diametrically opposite one another from the ghosts' origin, with varying ratios of distance and strength. Then we compute the percentage p1 of runs (out of a total of 45) that form a path to T1 rather than to T2. Figure 11 shows a plot of this probability as a function of the strength and distance ratios between the two targets. The dots represent experimental observations, between which other values are linear interpolations. Most of the plot is dominated by regions in which p_1 is either 1 or 0. The region within which both strength and distance play an active role in target selection is relatively narrow, and as both ratios grow, the difference in distance overwhelms the difference in strength.

Figure 11: Trade between target strength and distance

This plot is obtained with a single setting of pheromone parameters. It is possible to change the shape of this plot by adjusting the parameter settings. Appropriate techniques to automatically adjust these parameter settings to achieve the desired trade-off between distance and target strength is the subject of future research. Some preliminary experiments with evolutionary computing indicate that these techniques hold significant promise to address this issue [15].

Another important trade in understanding the behavior of ghost agents is between time and distance. When they are far from a target, ghost agents execute a random walk. Closer to the target, they can sense the target's pheromone field, and climb its gradient. One might expect that the number of steps required to reach a target would increase precipitously as the distance between a ghost's origin and its target grows. In fact, the transition is quite well behaved, as Figure 12 shows.


Figure 12: Time for first ghost to reach target

3.4 Visualization

Any system like synthetic pheromones that relies on emergent behavior poses a significant challenge for humans responsible for the system. The detailed behaviors of individual agents are driven by a wide range of environmental influences, and are impossible to predict in advance. Human users are willing to sacrifice detailed predictability to gain system-level robustness and adaptability, but they need to monitor system-level behavior for consistency with operational requirements, so that they can adjust the environmental influences impinging on the agents as needed.

A considerable portion of our effort in synthetic pheromones has been devoted to how to visualize them and their effects on agents for consumption by humans. The challenge is two-fold.

First, the system embraces a large number of entities of many different types, and can quickly overwhelm the user with detail. A typical run includes hundreds of avatars representing different classes of entities, thousands of ghosts, and dozens of flavors of pheromones, all spread over a graph of hundreds or thousands of place agents. We have explored a range of representational schemes for these entities, including shading of place agents to represent pheromone strengths and agent populations, discrete icons for entities of different types, and icon size as an indicator of strength or population. For experimental purposes, we have found simple shading mechanisms most effective, but these have not been validated against considerations of ease-of-use and clarity with actual human users, a necessary step in moving such systems into real-world deployment.

The second challenge in visualizing synthetic pheromone systems is how to provide a useful central view of an essentially distributed and decentralized system. In our current work, we "cheat," relying on the fact that our system runs on a single multitasking computer and driving visualization from a monolithic log file. However, even in constrained experimental situations, these logs are frequently on the order of 1Gbyte in size, and would be impossible to manage for a real situation. We are currently exploring mechanisms for using distributed visualization agents to gather information selectively from the distributed population of place and walker agents and using emergent mechanisms to digest this information for use by humans.

4. OPERATIONAL SCENARIOS

We have demonstrated the applicability of these mechanisms to military air operations in three increasingly sophisticated scenarios.

4.1 SEADy Storm

SEADy Storm [7] is a war game used to explore technologies for controlling air tasking orders. The battlespace is a hexagonal grid of sectors, each 50 km across Figure 13). Friendly (Blue) forces defend a region in the lower left against invading Red forces that occupy most of the field. Red's playing pieces include ground troops (GT's) that are trying to invade the Blue territory, and air defense units (AD's, surface-to-air missile launchers) that protect the GT's from Blue attack. Blue has bombers (BMB's) that try to stop the GT's before they reach the blue territory, and fighters tasked with suppressing enemy air defenses (SEAD's).

Each class of unit has a set of commands from which it periodically chooses. Ground-based units (GT and AD) choose a new command once every 12 hours, while air units (BMB and SEAD) choose once every five minutes. These times reflect the time it would take the resource to move across a sector. The commands fall into three categories (Table 2). GT cannot attack Blue forces, but can damage BMB's if they attack GT.


Figure 13: The SEADy Storm Game Board

Table 2: Unit Commands in SEADy Storm.

	Move	Attack	Wait
AD	Relocate	Fire (on any	Hide
		Blue aircraft)	Deceive
GT	Advance		Hide
SEAD	NewSectors	AttackAD	Rest
BMB	NewSectors	AttackAD	Rest
		AttackGT	

Blue can attack AD and GT when they are moving or attacking, and AD may attack any Blue forces that are not moving or waiting. Each unit has a strength that is reduced by combat. The strength of the battling units, together with nine outcome rules, determine the outcome of such engagements. Informally, the first five rules are:

1. Fatigue: The farther Blue flies, the weaker it gets.

2. Deception: Blue strength decreases for each AD in the same sector that is hiding.

3. Maintenance: Blue strength decreases if units do not rest on a regular basis.

4. Surprise: The effectiveness of an AD attack doubles the first shift after the unit does something other than attack.

5. Cover: BMB losses are greater if the BMB is not accompanied by enough SEAD.

Rules 6-9 specify the percentage losses in strength for the units engaged in a battle, on the basis of the command they are currently executing. For example, Rule 9, in full detail, states: "If BMB does "AttackGT" and GT does "Advance": a GT unit loses 10% for each BMB unit per shift; a BMB unit loses 2% per GT unit per shift."

The primary parameter explored in the experiments reported here is the proportion of SEAD in the Blue military, and of AD in the Red military. Each side began with a 100 units, each with unit strength, and 10%, 20%, 50%, 80%, or 90% of SEAD or AD. The uneven spacing reflects a basic statistical intuition that interesting behaviors tend to be concentrated toward the extremes of percentage-based parameters. In current military doctrine, 50% is an upper limit on both AD and SEAD. We explore higher values simply to characterize the behavioral space of our mechanisms.)

The central outcome is total Red strength in Blue territory at the end of the run (Figure 14). The landscape shows several interesting features, including

- a "valley" of Blue dominance for all Red ratios when Blue SEAD is between 50% and 80%, with slightly increasing Red success as the AD proportion increases;

- clear Red dominance for lower SEAD/BMB ratios, decreasing as SEAD increases;

- a surprising increase in Red success for the high SEAD and low AD levels.

Figure 15 compares the population of Red in Blue territory as a function of red and blue force composition for three different Blue control strategies. In the left plot, Blue does not use pheromones at all. The variations are due to the intrinsic dynamics of the combat, yielding a narrow valley up the center of the plot where Red's population is 3 or less (the criterion

Figure 14: Red Strength in Blue as a function of force composition

for Blue victory). When Blue uses pheromones to seek out Red targets and threats (center plot, shown in profile in Figure 14), the wider valley reflects improved Blue performance. In the right figure, when Blue uses pheromones to avoid threats and approach targets, the valley with the lowest Red population is about the same area but of a very different shape than in the previous case, but the next level of Red occupation (4-6) is much larger, showing a reduction in higher levels of Red occupation.

A detailed discussion of the dynamics of this scenario and interesting effects when we change the modeling formalism is available at [9].


Figure 15: Population of Red in Blue without pheromones (left), with pheromones attracted to both targets and threats (center), and attracted to targets but repelled by threats (right)

4.2 CyberStorm

At the next level of sophistication, we expand the range of unit types. Red now has armored and infantry battalions, air defense units, distinct headquarters types for regiments, air defense, and the entire corps, and fueling stations. Blue has three types of fighters and two types of bombers. The environment includes bridges and road crossings (which speed the movement of ground units that encounter them) and oil fields (which Red seeks to attack and Blue seeks to protect). Combat outcome is based on the percentage survival of the oil fields.

Using this enriched environment, we have explored a variety of issues around blue decision-making. In these experiments (as in SEADy Storm), Blue resources move directly in response to Red pheromones, without using ghosts. Our experiments show that reasonable numbers of Blue resources cannot sample the pheromone field

adequately to overcome the stochasticity inherent in the domain. As a result, outcomes vary widely with random seeds. These experiments demonstrated the need for ghost agents to sample the primary pheromone field at a statistically more significant level, and preprocess it for use by Blue avatars and the physical resources with which they are associated.

4.3 Super Cyber Storm

We exercised the ghost agents on a third model of the domain, which includes a significantly wider range of entity types, combat resolution on the basis of individual weapon type rather than unit type, more realistic dependencies among entities (for example, the effectiveness of Red air defense now depends on the status of other Red air defense units), and most importantly, a "pop-up" Red capability that lets us increase greatly the range of changes in Red's visibility as a scenario unfolds. This environment permits us to assess the effectiveness of ghost-based pheromones in dealing with pop-up threats.

First, we make all Red threats visible and stationary, and let the ghosts plan safe paths to the target for each of 181 offensive Blue missions against an entrenched Red force. We compute these paths using two different propagation parameters for Red threat pheromones, one that permits paths to fly relatively close to the threats, and another that keeps paths relatively far from the threats. Then we turn on Red movement and hiding behaviors, and compare the outcome of two sets of runs. In one set, Blue does not use ghosts or pheromones at all, but simply flies each mission on its precomputed path. This mode of operation corresponds to traditional pre-planned flight itineraries, except that our pre-planned paths, based on complete knowledge of Red's locations at the time of planning, are superior to those that could be constructed in a real conflict. In the other set of runs, Blue ignores precomputed paths and relies on ghosts to form paths for its missions dynamically. We assess the outcome of each run by the total remaining strength of Blue and Red assets at the end of the set of missions.

Figure 16 shows the medians over five runs of Red and Blue total unit strengths for three configurations. In "pathscript," each mission flies the path precomputed for it using a high Red propagation parameter, leaving a conservative margin around Red threats. In "pathscriptnarrow," Blue again flies precomputed paths, this time using paths computed with a lower Red propagation parameter, and permitting Blue to come closer to Red locations. These less conservative paths lead to increased combat between Blue aircraft and Red threats, and both Red and Blue losses increase compared with "pathscript." In "pathghost," Blue missions ignore precomputed paths and send out ghosts to compute their paths dynamically as the mission unfolds. In this mode of operation, Blue's losses are least, since it can now avoid pop-up Red threats. As a result, it can deliver more

Figure 16: Effect of Ghost Agents in Super Cyber Storm

weaponry to its assigned targets, increasing Red's losses in comparison with the other two scenarios. These differences would be even greater in an actual scenario where path planning for "pathscript" and "pathscriptnarrow" did not have include complete knowledge of all Red unit locations.

5. ARCHITECTURAL IMPLEMENTAION ISSUES

The overall architecture has been introduced in the previous sections. In this section we focus on to instantiate this architecture in a realistic C2 system. Figure 17 depicts the architecture of ADAPTIV and how it fits in with aspects of the overall C2 architecture.
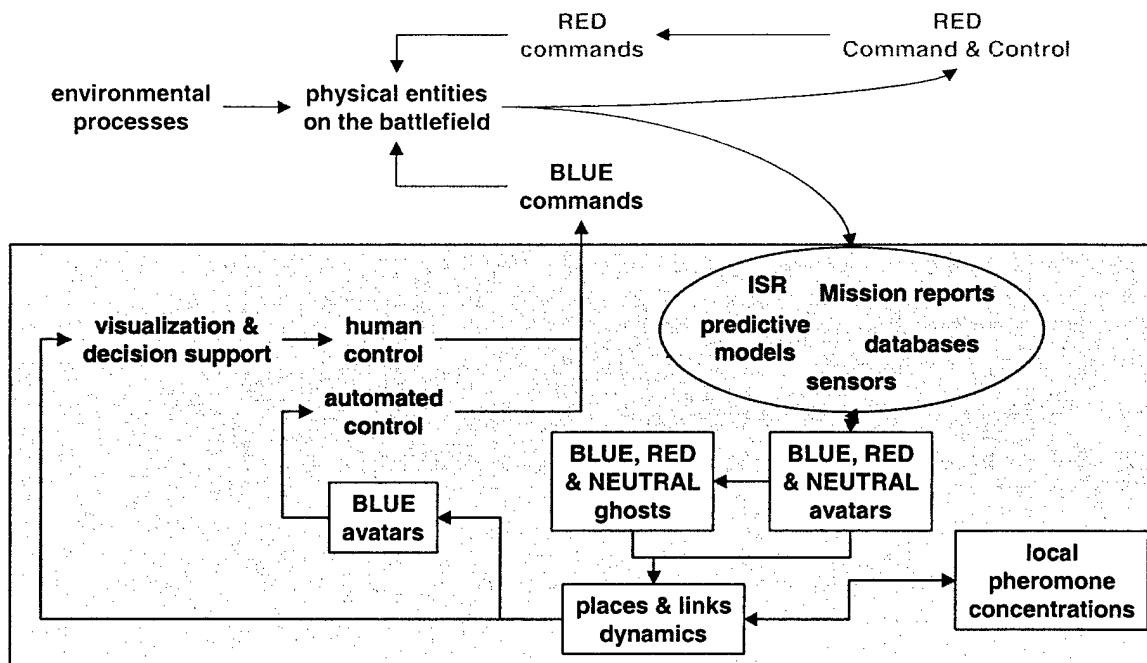
Figure 17: ADAPTIV C2 Architecture

The ADAPTIV system processes incoming observations, reports, and predictions. This information can include ISR, mission reports, predictive models, sensors, databases, or any other modality that can yield information about the battlespace. ADAPTIV's pheromone infrastructure provides a common integrating representation for, and dynamic maintenance of, this information. ADAPTIV's output consists of safe path formation, identification of alternative targets and potential actions to be taken by each avatar in response to changes in the environment. This output can be directed along two different channels (or a combination of the two):

- Direct control of Blue entities in the battlespace such as unmanned robotic vehicles.

- Input to human decision makers as a decision support tool.

Visualization can be used to view the emergent paths being formed by the ghost agents. It can show alternative targets that the ghosts may have identified giving decision makers an opportunity to re-direct missions enroute.

Red avatars are used to model and represent what is known about the enemy. The current status and location of the enemy units is derived from any sources of ISR available to the commander. When ISR indicates the possible presence of an enemy unit, an avatar is automatically created to model that unit. Its behavior can be based on a combination of historical models of that class of unit as well as learned behavior models obtained from observing a particular unit in action. Certainty of information is maintained both by the relative strength of the unit's deposit (with higher strength assigned to greater certainty) as well as the dispersion of the pheromone deposit (with a wider dispersion used to model less precise location information). Once a Red avatar is instantiated, it immediately begins to decrease its deposit rate according to a schedule determined by the unit type. Additional ISR re-confirming the unit's location and status serves to re-enforce the deposit amount. If no additional information on the unit is received, then over time the deposit amount will be reduced below a threshold and the avatar will be removed from the system. Thus old information is automatically removed from the system.

The ADAPTIV software agents must be instantiated in some software environment, integrated with legacy systems, and executed on a hardware infrastructure that includes the computers and communications support. Figure 18 shows how each of the ADAPTIV agents could be instantiated on a single computer platform called the *Host* and situated within a C2 architecture.
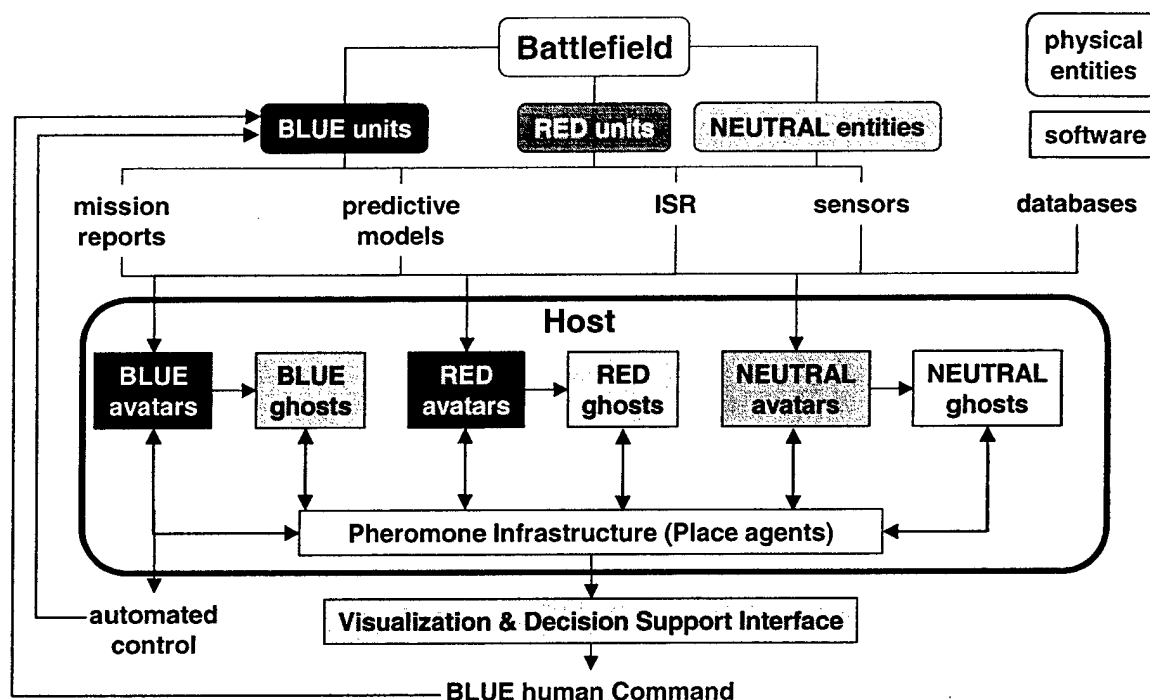
Figure 18: ADAPTIV on a Centralized Hardware Environment.

Because each agent in ADAPTIV is an independent process, it can be easily distributed across multiple computer platforms. Figure 19 depicts a notional example of how the agents in ADAPTIV could be distributed across several computers. In a traditional C2 architecture, different controllers responsible for different areas of the battlespace would house the place agents that represent those areas of the battlespace. As Blue and Red units move through the real battlespace, their avatars would move from place agent to place agent moving across computer boundaries as necessary. Information from the ADAPTIV agents would be visualized locally and centrally.

As computer and communications technology progresses the ADAPTIV architecture can migrate to even finer-grained, massively distributed computers. It may be possible to distribute small, camouflaged computers throughout the battlespace. These miniature hosts would contain a small computer, GPS, some local sensors, and short-range communication capability. Each host can represent a single place and would be responsible for maintaining the pheromone infrastructure for the portion of the battlespace it is responsible for (determined by its geographical location). Information received through local sensors and externally through its communications channels would be used to update its pheromone deposits. Manned and unmanned vehicles would communicate with the hosts to determine the pheromone concentrations locally to make decisions about what actions to take and what directions to move.

6. CONCLUSIONS

Synthetic pheromones are a powerful mechanism for controlling the movement of agents through space. They provide the elegance of potential field methods, with particular support for integrating diverse information sources, processing information in a completely distributed and decentralized environment, and coping with dynamic changes in the landscape. In exploring successively complex military scenarios, we have developed a toolkit of methods and mechanisms, including pheromone vocabularies that vary in both semantics and dynamics, mechanisms for incorporating agent momentum into movement decisions, ghost agents to preprocess the pheromone field and reduce stochasticity at the level of physical resources, and visualization mechanisms to enable human stakeholders to understand and monitor the emergent behavior of the system.
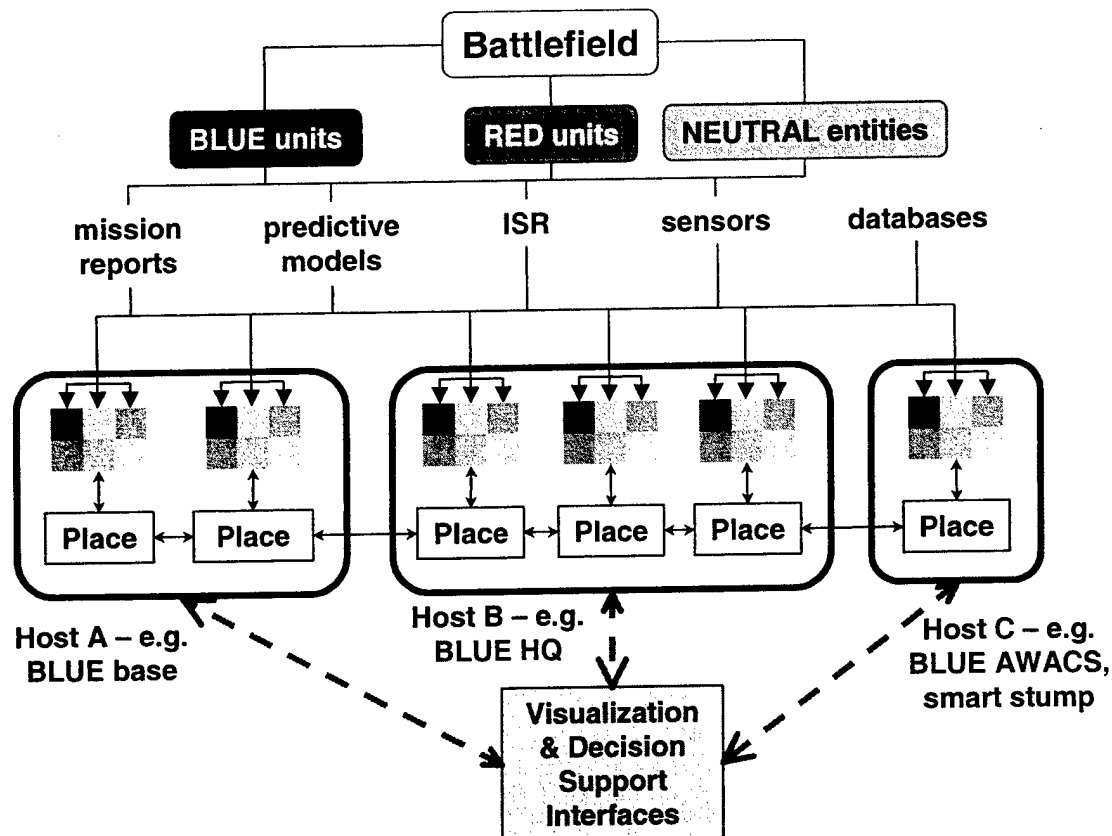
Figure 19: Distributed ADAPTIV Architecture.

7. ACKNOWLEDGMENTS

Realization of these concepts owes much to our research team, including E. Feibush, E. Greene, O. Gilmore, R. Matthews, and M. Nandula.

8. REFERENCES

[1]   E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. New York, Oxford University Press, 1999.

[2]   S. Brueckner. *Return from the Ant: Synthetic Ecosystems for Manufacturing Control*. Thesis at Humboldt University Berlin, Department of Computer Science, 2000.

[3]   S. Brueckner and H. V. D. Parunak. Multiple Pheromones for Improved Guidance. In *Proceedings of Symposium on Advanced Enterprise Control*, 2000.

[4]   A. Drogoul. When Ants Play Chess (Or Can Strategies Emerge from Tactical Behaviors? In *Proceedings of Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW '93)*, pages 13-27, Springer, 1995.

[5]   S. Goss, S. Aron, J. L. Deneubourg, and J. M. Pasteels. Self-organized Shortcuts in the Argentine Ant. *Naturwissenschaften*, 76:579-581, 1989.

[6]   D. Helbing, F. Schweitzer, J. Keltsch, and P. Molnár. Active Walker Model for the Formation of Human and Animal Trail Systems. Institute of Theoretical Physics, Stuttgart, Germany, 1998. URL http://xxx.lanl.gov/ps/cond-mat/9806097.

[7]   A. Kott. SEADy Storm. 2000. Word file, http://www.cgi.com/web2/govt/seadystorm.doc.

[8]   P. N. Kugler and M. T. Turvey. *Information, Natural Law, and the Self-Assembly of Rhythmic Movement*. Lawrence Erlbaum, 1987.

[9]   H. V. Parunak, S. Brueckner, J. Sauter, and R. Matthews. Distinguishing Environmental and Agent Dynamics: A Case Study in Abstraction and Alternative Modeling Technologies. In A. Omicini, R. Tolksdorf, and F. Zambonelli, Editors, *Engineering Societies in the Agents' World (ESAW'00)*, vol. LNAI 1972, *Lecture Notes in Artificial Intelligence*, pages 19-33. Springer, Berlin, 2000.

[10]   H. V. D. Parunak. 'Go to the Ant': Engineering Principles from Natural Agent Systems. *Annals of Operations Research*, 75:69-101, 1997.

[11]   H. V. D. Parunak. Adaptive control of Distributed Agents through Pheromone Techniques and Interactive Visualization. 2000. Web page, www.erim.org/cec/projects/adaptiv/.

[12]   H. V. D. Parunak and S. Brueckner. Ant-Like Missionaries and Cannibals: Synthetic Pheromones for Distributed Motion Control. In *Proceedings of Fourth International Conference on Autonomous Agents (Agents 2000)*, pages 467-474, 2000.

[13]   H. V. D. Parunak and S. Brueckner. Entropy and Self-Organization in Multi-Agent Systems. In *Proceedings of International Conference on Autonomous Agents*, pages (forthcoming), 2001.

[14]   E. Rimon and D. E. Kodischek. Exact Robot Navigation Using Artificial Potential Functions. *IEEE Transactions on Robotics and Automation*, 8(5 (October)):501-518, 1992.

[15]   J. A. Sauter, H. V. D. Parunak, S. Brueckner, R. Matthews, Tuning Synthetic Pheromones With Evolutionary Computing, In *Proceedings of Workshop on Evolutionary Computation and Multi-Agent Systems* (forthcoming), July 2001.

MISSION
OF
AFRL/INFORMATION DIRECTORATE (IF)

The advancement and application of Information Systems Science

and Technology to meet Air Force unique requirements for

Information Dominance and its transition to aerospace systems to

meet Air Force needs.